arXiv:1605.05402v1 [physics.soc-ph] 18 May 2016

Statistical mechanics of ontology based annotations

David C. Hoyle

Thorpe Informatics Ltd.,
Adamson House,
Towers Business Park,
Wilmslow Rd., Manchester,
M20 2YY, UK.

Andrew Brass

School of Computer Science,
University of Manchester,
Kilburn Building,
Oxford Rd., Manchester,
M13 9PL, UK.

Abstract

We present a statistical mechanical theory of the process of annotating an object with terms selected from an ontology. The term selection process is formulated as an ideal lattice gas model, but in a highly structured inhomogeneous field. The model enables us to explain patterns recently observed in real-world annotation data sets, in terms of the underlying graph structure of the ontology. By relating the external field strengths to the information content of each node in the ontology graph, the statistical mechanical model also allows us to propose a number of practical metrics for assessing the quality of both the ontology, and the annotations that arise from its use. Using the statistical mechanical formalism we also study an ensemble of ontologies of differing size and complexity; an analysis not readily performed using real data alone. Focusing on regular tree ontology graphs we uncover a rich set of scaling laws describing the growth in the optimal ontology size as the number of objects being annotated increases. In doing so we provide a further possible measure for assessment of ontologies.

Keywords: Information theory, Ontology, Zipf's law, Scaling law, Annotation

Email addresses: david.hoyle@thorpeinformatics.co.uk (David C. Hoyle), a.brass@manchester.ac.uk (Andrew Brass)

Published in *Physica A*, http://dx.doi.org/10.1016/j.physa.2015.09.020

©2016. This manuscript version is made available under the CC-BY-NC-ND 4.0 license http://creativecommons.org/licenses/by-nc-nd/4.0/

1. Introduction

With larger and more complex data sets becoming increasingly common, the annotation of data in order to semantically enrich it is a crucial task within data science [1, 2]. For example, in molecular biology a gene can be annotated by domain experts with terms, t, from a controlled vocabulary, thereby allowing other researchers to comprehend the function and role of that gene. Similarly, user tagging of information sources such as documents, photographs, or online content, provide additional meta-data and lead to emergent but uncontrolled vocabularies (often called folksonomies [3]). Terms within a vocabulary can be further organized in a hierarchical structure such as a taxonomy [1], in which terms closer to the root of the hierarchy are less specific than those further from the root. Hierarchical organization of vocabulary terms can also be used to specify richer semantic relationships between terms - richer than just indicating that one term is a simply a sub-type of another. These richer hierarchical semantic structures are generally called ontologies [4].

The annotations that result from an ontology or taxonomy can exhibit interesting patterns. For example, Kalankesh *et al.* [5] have shown that distributions of term frequencies, f_t, taken from Gene Ontology (GO) [6] annotations typically follow Zipf's law [7, 8]. Figure 1 shows a Zipf's law plot for annotations taken from the cellular component sub-ontology of GO. The schematic on the right-hand side of Fig.1 shows, for illustration, part of the ontology that was used to produce the annotation data set plotted on the left-hand side of Fig.1. Statistical mechanics provides us with a natural tool to understand these annotation patterns, by allowing us to develop a formalism that quantifies both the natural variations in the annotation process, and the ontology structure itself. Although structure-based measures of ontologies already exist [9], within this work we are quantifying the ontology structure from the perspective of the annotations that arise, rather than simply quantifying the ontology structure in isolation. The goals, and ultimately the benefits, of developing a statistical mechanics based formalism are both practical and theoretical.

1.1. Ontologies as information stores

Ontologies and taxonomies, whether formally constructed or emergent, represent a store of information. Organizing a hierarchical store of information requires effort to be expended to create an ordered structure. Work by Ferrer i Cancho *et al.*, within an information theoretic framework, has shown how heavy tailed and essentially hierarchical patterns of term usage can arise simply from a principle of minimizing the communication effort expended when using those terms [10, 11, 12]. Within this current paper we also use information theoretic ideas, but it is the process of transferring information from an ontology to an annotated object that we study, *i.e.*, *after* the hierarchical term structure has been determined or prescribed. We do so using an explicit statistical mechanical model that takes into the structure of the ontology. Whilst existing work within the literature has used a specific Hamiltonian to study patterns of word usage, that work has not *per se* been interested in the impact of any underlying

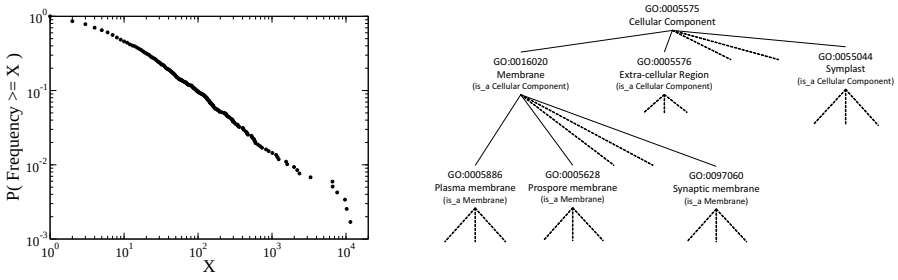

Figure 1: Left-hand plot shows a Zipf's law plot for GO gene annotations contained within Human GOA. Annotations have been taken from the cellular component sub-ontology. Terms near the root of the cellular component of GO are shown, for illustration, in the right-hand schematic - the dashed lines indicate the presence of further child terms.

prescribed structure in the vocabulary [13]. Similarly, novel work by Palla *et al.* [14] and Tibély *et al.* [15] has related tag usage patterns to ontology structure, but focused on an in depth study of observed tag patterns, rather than taking a Hamiltonian model based approach.

The remainder of this paper is organized as follows - in Section 2 we express the annotation process as an ideal lattice gas model in an inhomogeneous field. In Section 3.1 we use the lattice gas model to understand the term frequency patterns seen by Kalankesh *et al.* [5], and we identify, LD_t, the number of leaf descendants of a node t, as the key quantity controlling the expected term usage frequencies. In Section 4 we derive the most likely natural form for the inhomogeneous field strength, thereby giving rise to a local measure of the ontology. This natural form for the inhomogeneous field also allows us, in Section 5, to construct an ensemble of ontologies of differing complexity. By restricting the ensemble to the class of regular trees we reveal in Section 5.2 a set of transitions in the optimal tree size, and associated scaling laws, as the number of objects being annotated is increased. Finally in Section 6 we discuss a number of possible extensions of the statistical mechanical approach to quantifying ontology structures.

2. Statistical mechanical theory of the annotation process

We consider an ontology to be represented by a rooted Directed Acyclic Graph (DAG) [16, 17]. An example DAG, in this case a tree, is shown in Figure 2. Real-world ontologies are typically not pure trees, and we use a tree structure simply for illustrative purposes. The formalism we develop in this section will be equally applicable to any valid DAG structure. Associated with each node of the DAG is a particular term, and we use node and term interchangeably.

The directed edges between nodes of the DAG indicates that the child node conveys a possibly more specific meaning than the parent node, or represents a more specific subset of objects. For such an increase in specificity to occur the ontology topology must be ultimately be tree-like. Clearly, if a particular

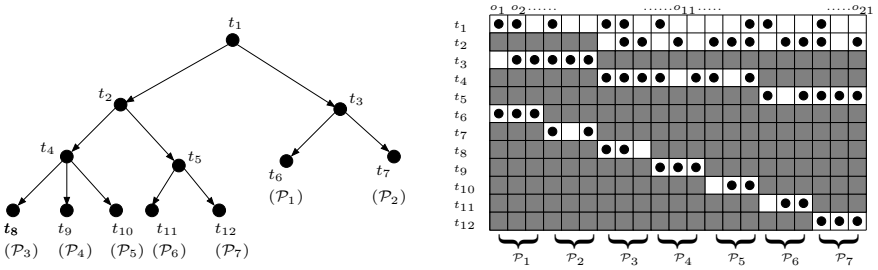

Figure 2: Left hand schematic shows an example (tree) DAG representing an ontology. The nodes represent terms, denoted by $t_1 - t_{12}$, which may be selected to annotate objects that belong to classes represented by the leaves of the DAG. Each leaf corresponds to a unique set of paths from the root, and so we denote the classes by $\mathcal{P}_1 - \mathcal{P}_7$. The directed edges of the DAG indicate semantic relationships between the terms. The right-hand side shows a lattice-gas schematic of an example annotation, *i.e.* a selection of terms appropriate to 21 objects, $o_1, o_2, \ldots, o_{21}$, that have been assigned to the 7 classes $\mathcal{P}_1 - \mathcal{P}_7$ (3 objects per class). A cell occupied by a particle in the lattice-gas schematic indicates the term has been selected to annotate that object, whilst a filled (grey) cell indicates a cell that cannot be occupied as a consequence of the ontology structure.

term is applicable to an object then so are the terms associated with any of the ancestor nodes.

At a leaf, we have a term that is applicable to a small subset of highly defined objects, possibly even a single object. Each leaf of the DAG determines a unique path or set of paths, denoted by $\mathcal{P}$, consisting of all terms that can be traversed in moving from the root to the leaf. Thus we use leaf, class and $\mathcal{P}$ interchangeably and use $t \in \mathcal{P}$ to denote a term on the unique set of paths $\mathcal{P}$. The prior probability of a particular class $\mathcal{P}$ is $\pi_{\mathcal{P}}$, and is just the proportion of all objects within the class $\mathcal{P}$. If, within a data set we have a total of $|\mathcal{O}|$ objects and $|\mathcal{O}_{\mathcal{P}}|$ objects from class $\mathcal{P}$, then we can construct a simple estimator, $\hat{\pi}_{\mathcal{P}}$, of $\pi_{\mathcal{P}}$ via,

$$\hat{\pi}_{\mathcal{P}} \;=\; \frac{|\mathcal{O}_{\mathcal{P}}|}{|\mathcal{O}|} \quad . \tag{1}$$

Clearly, we expect $\hat{\pi}_{\mathcal{P}} \to \pi_{\mathcal{P}}$ as $|\mathcal{O}| \to \infty$. We then use $\mathcal{N}_t$ to denote the expected number of classes within the data set to which a term t belongs, *i.e.*, $\mathcal{N}_t$ is simply the number of leaves which are descendants of term t in the ontology, weighted by the prior $\pi_{\mathcal{P}}$. Formally we define,

$$\mathcal{N}_t \;=\; \sum_{\mathcal{P} \ni t} \pi_{\mathcal{P}} \quad , \tag{2}$$

and denote by $\widehat{\mathcal{N}}_t$ its equivalent defined using the estimators $\hat{\pi}_{\mathcal{P}}$, *i.e.*,

$$\widehat{\mathcal{N}}_t \;=\; \sum_{\mathcal{P} \ni t} \hat{\pi}_{\mathcal{P}} \quad . \tag{3}$$

The value of $\mathcal{N}_t$ is essentially the probability, $P(t)$, that term t is relevant to an object, irrespective of any further details of the object. Thus $\mathcal{N}_t$ gives a probabilistic measure of the specificity of term t that reflects the structure of the ontology, and expert knowledge and information encoded within it. We therefore consider $\mathcal{N}_t$ to measure the intrinsic information content of term t. Specifically, we define the information content of term t as $-\log P(t) = -\log \mathcal{N}_t$ [18].

An object belonging to the class represented by a leaf $\mathcal{P}$ can be annotated with any of the terms t belonging to $\mathcal{P}$. We use $n_{ot} \in \{0, 1\}$ to denote whether term t is used to annotate object o. The discrete variables n_{ot} are lattice gas occupancy numbers (or equivalently Ising spins). Consequently, we formulate the probability of a particular choice of annotations as a lattice-gas model. An example lattice-gas configuration corresponding to an example annotation is shown by the schematic on the right-hand side of Fig.2. The lattice-gas Hamiltonian is relatively simple and is given by,

$$H \;=\; \sum_{o,t} n_{ot}(v_{ot} - \mu) \;\; . \tag{4}$$

Here v_{ot} represents the local field acting upon a particle at term t for object o, and determines how likely it is that term t will be selected when annotating object o. Ideally no mis-annotation occurs, *i.e.* only terms appropriate to each object are selected, and there would be the same probability of selecting term t provided the paths $\mathcal{P}$ to which the object belongs contains t. Thus, we set,

$$v_{ot} \;=\; \begin{cases} v_t & t \in \mathcal{P} \text{ and } o \in \mathcal{P} \;\; , \\ h & \text{otherwise} \;\; . \end{cases} \tag{5}$$

The value of h determines the global background level of mis-annotation that may occur. It is possible to consider a more structured form for v_{ot}, *e.g.*, to allow for greater likelihood of mis-annotation when a term t is close to, but not in, the paths $\mathcal{P}$ to which the object o belongs. However, for the remainder of this paper we consider the more simplified form for v_{ot} given in Eq.(5) above and also only consider the scenario where no mis-annotation occurs. Consequently, we set $h = \infty$. Therefore, for each object, particular terms are forbidden if the term is not on any of the paths from the root node to the leaf node associated with the object. This is also illustrated in Fig.2. In practice, when analysing real data sets, it may be required to consider a finite value of h, to capture the mis-annotations that will inevitably occur.

Within Eq.(4) we have also included a chemical potential, μ, which acts as a global field acting on all lattice cells. Equivalently, the fugacity $z = \exp(\beta\mu)$ is essentially the global prior probability that a term will be selected by an annotator, and controls or limits the average number, $\bar{n}$, of terms per object selected by an annotator. We have used $\beta = 1/T$ to denote the inverse temperature. The temperature T does not at this stage have an explicit physical interpretation, other than to represent a parameter which controls the expectation value of the Hamiltonian in Eq.(4). However, in Section 4 we propose an intrinsic

form for the external field v_t in terms of the information content of the term t. With this intrinsic form for v_t the temperature T then controls the average information retrieved from the ontology when an object is annotated. Clearly the temperature T will also determine the level of variation seen in annotations from otherwise identical annotators when annotating the same set of objects.

Since we only have non-interacting particles, the probability p_t, of term t being used to annotate an object for which it is valid, is easily calculated as,

$$p_t \;=\; \langle n_{ot} \rangle \;=\; \frac{e^{-\beta(v_t - \mu)}}{1 + e^{-\beta(v_t - \mu)}} \;\; . \tag{6}$$

Here $\langle n_{ot} \rangle$ denotes the expectation (over annotations) of n_{ot}. Likewise, with the simplification in Eq.(5) for v_{ot} the partition function is straight-forward to evaluate, and for a particular set of objects is given by,

$$\log Z \;=\; |\mathcal{O}| \sum_t \widehat{\mathcal{N}}_t \log \left(1 + e^{-\beta(v_t - \mu)} \right) \;\; . \tag{7}$$

Consequently, the expected value of $\log Z$, averaged over all possible data sets, $\mathcal{O}$, of size $|\mathcal{O}|$, is,

$$\langle \log Z \rangle_{\mathcal{O}} \;=\; |\mathcal{O}| \sum_t \mathcal{N}_t \log \left(1 + e^{-\beta(v_t - \mu)} \right) \;\; . \tag{8}$$

Likewise, the expected frequency, C_t, of term t within the annotation data set is given by,

$$C_t \;=\; \sum_{o \ni \mathcal{P} \ni t} \langle n_{ot} \rangle \;=\; -\beta^{-1} \frac{\partial \log Z}{\partial v_t} \;=\; |\mathcal{O}| \frac{e^{-\beta(v_t - \mu)}}{1 + e^{-\beta(v_t - \mu)}} \widehat{\mathcal{N}}_t \;=\; |\mathcal{O}| p_t \widehat{\mathcal{N}}_t \;\; , \tag{9}$$

with the expectation over all data sets of size $|\mathcal{O}|$ being,

$$\langle C_t \rangle_{\mathcal{O}} \;=\; |\mathcal{O}| \frac{e^{-\beta(v_t - \mu)}}{1 + e^{-\beta(v_t - \mu)}} \mathcal{N}_t \;=\; |\mathcal{O}| p_t \mathcal{N}_t \;\; . \tag{10}$$

In all of the expressions above for the partition function Z we do not consider objects within the same class $\mathcal{P}$ to be interchangeable. It is likely that objects within the class $\mathcal{P}$ are unique and distinguishable, and simply possess a common characteristic rather than being replicates of each other.

Finally, we note that for any real data set of term frequencies f_t, Eq.(10) provides a simple means of constructing an estimate of the field strength operating on node t, by equating the expectation value for $\langle C_t \rangle_{\mathcal{O}}$ given in Eq.(10) with f_t. This gives us an estimate $\hat{v}_t$ for v_t given by,

$$\hat{v}_t \;=\; \beta^{-1} \left[\log z \;+\; \log \left(\frac{|\mathcal{O}| \mathcal{N}_t}{f_t} - 1 \right) \right] \;\; . \tag{11}$$

3. Replicating the patterns of real annotation data sets

3.1. The relation between $\mathcal{N}_t$ and leaf descendants

Clearly, a first test of the lattice-gas model of the annotation process is whether it can replicate or explain the broad patterns seen in real annotation data sets, such as those seen by Kalankesh *et al.* [5]. From Eq.(10) we can see that $\langle C_t \rangle_{\mathcal{O}}$ depends upon two factors; $\mathcal{N}_t$ and p_t. Due to the typically hierarchical nature of the ontology topology we would expect $\mathcal{N}_t$ to have a broad distribution of values, and potentially to dominate the distribution of $\langle C_t \rangle_{\mathcal{O}}$. It is transparent that for class probabilities that are uniform across all possible classes then the value of $\mathcal{N}_t$ is simply proportional to the number of leaf descendants, LD_t, that can be reached from term t. For small deviations away from a uniform class probabilities we would still expect $\mathcal{N}_t$ and LD_t to be approximately proportional. Larger deviations away from uniform class probabilities effectively represent a pruning of the DAG into a smaller topology, with leaf nodes that correspond to small class probabilities being effectively eliminated. Consequently, for any DAG topology we still consider LD_t to give useful insight into the expected term frequency $\langle C_t \rangle_{\mathcal{O}}$, and in the next section we assess the likely distribution for LD_t for tree-like ontologies.

3.2. The distribution of the leaf descendant count LD_t.

For irregular-trees precise evaluation of LD_t given local information about the node t can be framed in terms of Galton-Watson processes [19, 20], though there are few usefully applicable closed-form results available to us. In contrast, single parent regular trees are easier to study. For a single-parent regular tree with b children per node we can label the layers of the tree from $l = 0$ (at the root) to $l = L$ (at the leaves). The total number of nodes in the tree is $(b^{L+1} - 1)/(b - 1)$. The number of nodes in layer l is simply b^l, whilst the number of leaf descendants for each node in layer l is b^{L-l}, *i.e.*, within each layer there is a simple reciprocal relationship between node counts and leaf descendant counts, suggesting a Zipf's behaviour for the distribution of LD_t. More formally, for node t in layer l, with leaf descendant count $LD_t = b^{L-l}$, it is then a simple matter to find,

$$\text{Fraction of nodes with leaf descendant count} \geq LD_t \;=\; \frac{b-1}{b^{L+1}-1} \sum_{k=0}^{l} b^k$$

$$= \quad \frac{b^{l+1}-1}{b^{L+1}-1} \;=\; \frac{b^{L+1}}{b^{L+1}-1} \left[\frac{1}{LD_t} \;-\; \frac{1}{b^{L+1}} \right] \;\simeq\; \frac{1}{LD_t} \;, \text{ as } L \to \infty \,. \qquad (12)$$

Consequently, for large regular trees we will have $P(LD_t \geq X) \;\simeq X^{-1}$, *i.e.*, a Zipf's law form. Overall, Eq.(12) suggests the origin of the Zipf's law behaviour observed by Kalankesh *et al.* [5] may be, in part, a consequence of the Zipf's law like behaviour of the distribution of LD_t. However, whilst the structure of a regular tree will lead, on average, to a classical Zipf's law like behaviour, the degeneracy in leaf descendant count for nodes within the same layer leads to a

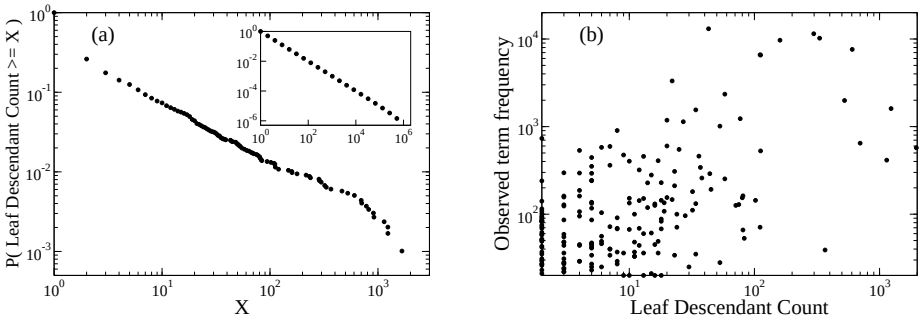

Figure 3: Plots of the distributions of leaf descendant counts for a real-world ontology and a regular tree. Plot (a) shows a plot of the log of the cumulative probability of LD_t for all nodes with children within the GO cellular component ontology. The inset within plot (a) shows a plot of the log of the cumulative probability of LD_t for all nodes with children within a bifurcating ($b = 2$) regular tree of depth $L = 20$. Plot (b) shows a log-log plot of observed term frequencies, f_t, against number of leaf descendants, LD_t, for the real data set shown in Fig.1.

very evenly spaced behaviour in any plot of the cumulative probability. This can be seen in the inset of Figure 3a which shows the cumulative probability distribution (on a logarithmic scale) of leaf descendant counts for a regular tree with $b = 2$ and depth $L = 20$.

For a real ontology, where we will have significant variation in the number of children and parents at the scale of individual nodes, we would expect to see more continuous cumulative probability plots for LD_t. We have confirmed with further simulations of irregular tree-like ontologies (results not shown) that, typically, continuous Zipf's law like behaviour in LD_t results when one has local variation in the node characteristics, *e.g.*, variation in the average number of children per node or in the probability that a node has children. Similarly, the main part of Fig.3a shows the cumulative probability distribution plot for the leaf descendant count of the cellular component GO ontology. This is the ontology used for annotating the data set shown in Fig.1. The intrinsic power-law like behaviour is clearly evident, and so it perhaps unsurprising that we should observe power-law like behaviour in annotations based upon this ontology. As postulated, and in contrast to the distribution shown for the regular tree, this real-world ontology shows a more continuous spread of leaf descendant counts. However, as with the regular tree, the effective exponent of the approximate power-law form for the cellular component GO ontology is close to -1.

3.3. The external field v_t: intrinsic and specific components

We have argued that term frequencies in real annotation data sets will be broadly determined by LD_t. Figure 3b shows a log-log plot for the observed term

frequencies, f_t, of the data set shown in Fig.1, plotted against the number of leaf descendants within the ontology from which the annotation terms were drawn. The broad correspondence between observed log of the term frequencies and the log of the number of leaf descendants is clear, and is statistically significant (estimate of Pearson correlation $\hat{\rho} = 0.454,\ 95\%CI = [0.408, 0.498]$, $p < 10^{-8}$ for null hypothesis of $\rho = 0$).

However, we would not expect an exact correspondence with the real observed term frequencies, f_t, on the basis of $\mathcal{N}_t$ alone. The precise value of $\langle C_t \rangle_{\mathcal{O}}$ is determined by two factors; $\mathcal{N}_t$ and p_t, with p_t determined by v_t. We might expect v_t itself to have an intrinsic component determined by the ontology topology, but the scatter in Fig.3b suggests that we should decompose v_t into an intrinsic contribution, $v_t^{(0)}$, and a specific residual contribution, Δv_t. That is we write,

$$v_t \ = \ v_t^{(0)} \ + \ \Delta v_t \tag{13}$$

Once a form for $v_t^{(0)}$ has been specified and an estimate for v_t has been obtained from observed term frequencies f_t via Eq.(11), then an estimate for Δv_t can be obtained. Overall, the values of Δv_t reveal any biases towards particular terms beyond that expected on the basis of the ontology topology alone. Significant values of Δv_t can potentially indicate regions of the ontology which are not matching the requirements of the annotators, and where the ontology can potentially be improved or needs modification. Equally, having a form for $v_t^{(0)}$ allows us to extend the theoretical statistical mechanical formalism by providing an explicit form for p_t in the absence of any bias in v_t. Therefore, in the remainder of the paper we derive an appropriate form for $v_t^{(0)}$ and use it to study the consequences of increasing ontology complexity.

4. Determination of the intrinsic field strength $v_t^{(0)}$

To proceed we note that extracting or transferring information from the ontology to an object during the annotation process requires effort. Selecting more specific terms, far from the root, requires greater information about the object to be known by, or inferred by the annotator, and so we equate effort expended by an annotator with the intrinsic information content of the terms selected. A natural choice for $v_t^{(0)}$ would therefore be to set $v_t^{(0)}$ proportional to the information content of term t, *i.e.*, set $v_t^{(0)} \propto -\log \mathcal{N}_t$. This choice corresponds to the maximum entropy solution for the term probabilities $\{p_t\}_t$ under the constraint of capturing a given amount of information from the ontology using a fixed expected number of terms. To see this we first calculate the entropy, $S[\{p_t\}]$ for for a given set of term probabilities, $\{p_t\}$. From Eq.(6) and Eq.(7) the entropy $S[\{p_t\}]$ is obtained as,

$$
\begin{aligned}
S[\{p_t\}] &= \langle \log Z \rangle_{\mathcal{O}} [\{v_t^{(0)}\}_t] \; + \; \beta \sum_{o,t} \langle \langle n_{ot} \rangle \rangle_{\mathcal{O}} (v_t^{(0)} - \mu) \\
&= -|\mathcal{O}| \sum_t \mathcal{N}_t \left[p_t \log p_t \; + \; (1-p_t) \log(1-p_t) \right] \; . \qquad (14)
\end{aligned}
$$

In calculating the entropy in Eq.(14) we have clearly averaged over data sets of size $|\mathcal{O}|$, as we wish to eliminate the effects of data set to data set variation and focus solely on the effect of data set size. If we wish to obtain the set of term probabilities $\{p_t\}$ which maximize the entropy subject to capturing a given amount of the stored information in the ontology and a given expected total number of selected terms, then we maximize,

$$
\begin{aligned}
& -\lambda \sum_{o,t} \langle \langle n_{ot} \rangle \log \mathcal{N}_t \rangle_{\mathcal{O}} \; + \; \phi \sum_{o,t} \langle \langle n_{ot} \rangle \rangle_{\mathcal{O}} \; + \; S\left[\{p_t\}\right] \\
= \; & -|\mathcal{O}| \sum_t \mathcal{N}_t p_t (\lambda \log \mathcal{N}_t - \phi) \; - \; |\mathcal{O}| \sum_t \mathcal{N}_t \left(p_t \log p_t \; + \; (1-p_t) \log(1-p_t) \right) \; . \qquad (15)
\end{aligned}
$$

Here λ and ϕ are Lagrange multipliers. The maximum entropy term probabilities, p_t^*, are simply the set of term probabilities that maximize the expression in Eq.(15), and so we find,

$$
p_t^* \; = \; \frac{e^{-(\lambda \log \mathcal{N}_t - \phi)}}{1 + e^{-(\lambda \log \mathcal{N}_t - \phi)}} \; . \qquad (16)
$$

Comparing Eq.(6) and Eq.(16) we immediately see that this maximum entropy approach corresponds to setting,

$$
\beta v_t^{(0)} \; = \; \lambda \log \mathcal{N}_t \quad , \quad \beta \mu \; = \; \phi \; . \qquad (17)
$$

With $\beta > 0$ we impose the relation (without loss of generality) $\beta = |\lambda|$, giving $v_t^{(0)} = -\log \mathcal{N}_t$ when $\lambda < 0$ and $v_t^{(0)} = \log \mathcal{N}_t$ when $\lambda > 0$. For $\lambda < 0$ terms with lower information content (larger $\mathcal{N}_t$), that are closer to the root, will be preferentially selected during annotation, whilst for $\lambda > 0$ terms with smaller values of $\mathcal{N}_t$, that are closer to the leaves of the ontology are preferentially selected.

Substituting the expression for p_t^* from Eq.(16) into Eq.(9) we get the corresponding maximum entropy expected term frequencies as,

$$
\langle C_t^* \rangle_{\mathcal{O}} \; = \; |\mathcal{O}| \frac{z \mathcal{N}_t^{1-\lambda}}{1 + z \mathcal{N}_t^{-\lambda}} \; . \qquad (18)
$$

Likewise, for the optimal term probabilities $\{p_t^*\}$, the free energy takes a value F^* given by,

$$
\beta F^* \; = \; \beta F[\{p_t^*\}_t, \{\beta v_t = v_t^{(0)}\}_t] \; = \; -|\mathcal{O}| \sum_t \mathcal{N}_t \log(1 + z \mathcal{N}_t^{-\lambda}) \; . \qquad (19)
$$

Having obtained a suitable form for $v_t^{(0)}$ it is a simple matter to determine whether, for any real data set, there is a non-zero specific field component Δv_t. This can be done by testing if the observed frequency f_t is statistically significantly different from $\langle C_t^* \rangle_{\mathcal{O}}$ given by Eq.(18). If f_t is significantly different to $\langle C_t^* \rangle_{\mathcal{O}}$, then we can then combine Eq.(11), Eq.(13), and Eq.(17) to construct an estimate $\widehat{\Delta v}_t = \hat{v}_t - v_t^{(0)}$ for Δv_t. Thus, we have,

$$\widehat{\Delta v}_t = \begin{cases} \frac{1}{|\lambda|}\left[\log\left(z\mathcal{N}_t^{-\lambda}\right) + \log\left(\frac{|\mathcal{O}|\mathcal{N}_t}{f_t} - 1\right)\right] & \text{if } f_t - \langle C_t^* \rangle_{\mathcal{O}} \text{ significant} \\ 0 & \text{otherwise .} \end{cases} \tag{20}$$

Realistically, for a large ontology we would only expect a small proportion of all the terms to be used when annotating an object, and so from Eq.(18) we would expect $\langle C_t^* \rangle_{\mathcal{O}}/|\mathcal{O}| \ll 1$. This suggests that $z\mathcal{N}_t \ll 1$ and that in general $\langle C_t^* \rangle_{\mathcal{O}} \sim \mathcal{N}_t^{1-\lambda}$. This gives us a simple mechanism for estimating appropriate values of λ for real data sets. This also tells us that, with $\mathcal{N}_t$ from tree-like ontologies expected to display a Zipf's law behaviour with an exponent close to -1, then we expect to see a Zipf's law like behaviour in the observed term frequencies f_t with exponent $-(1-\lambda)$. As the exponents for f_t observed by Kalankesh *et al.* [5] are close to -1, this suggests that the effective values of λ for real data sets are small in magnitude and potentially either side of $\lambda = 0$.

Finally, we note that as $\lambda = 0$ represents the boundary at which terms closer to the root are preferentially selected, an effective value of $\lambda < 0$ within a real data set may suggest that annotators do not have a strong desire to make use of the additional nodes provided by larger complex ontologies. Within our statistical mechanical framework we are able to explore whether there is a 'optimal' or preferred size to the ontology that annotators wish to use, and also how this 'optimal' size might scale with the number of objects being annotated. This we do in the next section.

5. Statistical mechanics of an ensemble of ontologies

Having set up a statistical mechanical model of the annotation process, we are in a position to perform more theoretical analyses and experiments that, whilst not being realisable, still provide valuable insight into the performance of annotators and the underlying ontology being used. For example, a key practical question we wish to address is to determine whether a real-world ontology is fit for purpose, or whether the ontology needs shrinking or expanding. To answer this question in real terms would require providing annotators with a range of ontology graphs and observing which ontology is used most frequently. Such an experiment is not readily performed in real terms. However, within the statistical mechanical framework we can address this question by extending the previous ensemble to one consisting of a large collection of annotators who annotate the same set of objects, but who can select the ontology to be used from a range of ontologies of differing complexity. Consequently, we will see variation

in both the choice of topology $\mathcal{H}$ being used, and the annotation configuration $\{n_{ot}\}$ being selected. Considering this ensemble allows us to easily determine the most preferred, and hence most appropriate, ontology. This optimal ontology will depend upon a number of parameters, such as the number of objects being annotated, $|\mathcal{O}|$, and annotator characteristics such as the average number of terms used per object and the average effort/information the annotator is willing to expend per object. The latter two characteristics are controlled by the parameters z and λ. Therefore, the remainder of Section 5 is focused on elucidating how the preferred, or optimal, ontology varies with the parameters $|\mathcal{O}|, z$ and λ. To determine the variation with respect to $|\mathcal{O}|, z$ and λ requires breaking the analysis down into a number of smaller, but still quite involved, steps. Firstly, in Section 5.1 we construct a suitable statistical mechanical potential, $\beta\Omega^*$, from which the optimal ontology topology $\mathcal{H}$ can be determined. In section 5.2 we restrict $\mathcal{H}$ to the class of regular trees to make analysis of $\beta\Omega^*$ tractable. In Section 5.3 we further simplify the analysis by developing a closed form approximation to $\beta\Omega^*$ for regular trees, and examining the asymptotic behaviour of the approximation as the number of objects $|\mathcal{O}|$ increases. In Section 5.4 we confirm the results of the asymptotic analysis on regular trees using numerical simulations. Finally, in Section 5.5 we use the insights gained from the analysis of $\beta\Omega^*$ to understand what are the most likely observed values for the parameter λ.

5.1. Quantifying ontology complexity and determination of the optimal ontology

To extend the ensemble to one that consists of ontologies of varying complexity we must first introduce a measure to quantifying the intrinsic complexity of each ontology. Continuing the information theoretic approach, we measure the complexity of an ontology by its total intrinsic information content $-\sum_t \log \mathcal{N}_t$, and so introduce an additional control variable, $\omega > 0$, to set the average intrinsic information content of a ontology node within this ensemble. The partition function for this ensemble becomes,

$$\mathcal{Z} \;=\; \sum_{\mathcal{H}} \exp\left(\omega \sum_t \log \mathcal{N}_t\right) Z(\mathcal{H}) \;\;, \tag{21}$$

where $Z(\mathcal{H})$ is the partition function on a fixed DAG and $\log Z(\mathcal{H})$ is given by $-\beta F^*$ in Eq.(19). Larger values of ω will favour smaller values for the average intrinsic information per node, and so will favour smaller, lower complexity ontologies. Just as the partition function has been modified, then similarly, the appropriate potential for this ensemble is obtained by adding $-\omega \sum_t \log \mathcal{N}_t$ to the optimal value of the free energy, βF^*. This gives a potential,

$$\beta\Omega^* \;=\; -|\mathcal{O}| \sum_t \mathcal{N}_t \log(1 + z\mathcal{N}_t^{-\lambda}) \;-\; \omega \sum_t \log \mathcal{N}_t \;\;. \tag{22}$$

The most frequently selected, or optimal, ontology is that which minimizes $\beta\Omega^*$. The term $-\omega \sum_t \log \mathcal{N}_t$ therefore serves an important function by regulating

the ontology complexity, with more complex topologies being penalized more heavily.

We can see that the minimum value $\beta\Omega^*(\{\mathcal{N}_t\}_t, z, \omega)$ changes on varying either the topology $\mathcal{H}$ or the total number of objects, $|\mathcal{O}|$, to which the ontology applies. Consequently, we have the possibility of a change in optimal DAG topology as the number of objects, $|\mathcal{O}|$, being annotated increases. As previously stated, this possibility of a change in the optimal DAG topology is something we wish to investigate further, though to do so over a range of general DAG topologies is unlikely to be tractable. However, as we have already argued for an ontology to be effective it will have a tree-like topology. This then highlights the importance of understanding the behaviour of $\beta\Omega^*$ for tree-like ontologies. Even so, characterizing the behaviour of $\beta\Omega^*$ for an arbitrary tree-like ontology is still likely to be a difficult task. Therefore, to gain further insight into the behaviour of $\beta\Omega^*$ for tree-like ontologies we restrict our further analysis of $\beta\Omega^*$ to regular trees.

5.2. Optimal ontology size for regular trees

For a regular tree with multi-furcating nodes, *i.e.*, each parent having b children, we can label nodes according to which layer, l, the node is in. The root node is in layer $l = 0$, whilst leaf nodes are in the final layer L. Nodes in layer l have b^{L-l} leaf descendants and there are b^l nodes in layer l. Consequently, the total number of leaves is b^L. For simplicity we will take the class distribution $\{\pi_{\mathcal{P}}\}_{\mathcal{P}}$ to be uniform across the leaf nodes, *i.e.*, $\pi_{\mathcal{P}} = b^{-L}$ for each of the classes corresponding to a leaf node. From these relations we can evaluate $\mathcal{N}_t = b^{-l}$, for a term corresponding to a node in layer l. It is then a simple matter to re-write $\beta\Omega^*$ as,

$$\beta\Omega^* \;=\; -|\mathcal{O}| \sum_{l=0}^{L} \log\left(1 + z b^{l\lambda}\right) \;+\; \omega \log b \sum_{l=0}^{L} l b^l \quad . \tag{23}$$

The average number, $\bar{n}$, of terms used per annotated object is determined by the fugacity z via the relation,

$$\bar{n} \;=\; -\frac{z}{|\mathcal{O}|}\frac{\partial}{\partial z}\beta\Omega^* \;=\; \sum_{l=0}^{L} \frac{z b^{l\lambda}}{1 + z b^{l\lambda}} \quad . \tag{24}$$

Thus, if we wish to attain a specified value $\bar{n}$ we simply solve Eq.(24) for the required value z. Similarly, the average information retrieved per object, I, is determined via the relation,

$$I \;=\; -\frac{1}{|\mathcal{O}|}\frac{\partial}{\partial \lambda}\beta\Omega^* \;=\; \log b \sum_{l=0}^{L} \frac{z l b^{l\lambda}}{1 + z b^{l\lambda}} \quad . \tag{25}$$

The simple form for the $\beta\Omega^*$ in Eq.(23) also allows us to analytically determine the tree size L that is optimal for annotating a given number of objects $|\mathcal{O}|$. The optimal tree size is that which minimizes $\beta\Omega^*$. Due to the discrete

nature of L, any growth we observe in the optimal ontology size will occur via a series of transitions. Naively, we would expect the optimal ontology size, L_{opt}, to increase as the number of annotated objects is increased. That is, we would *a priori* expect $L_{opt} \to \infty$ as $|\mathcal{O}| \to \infty$. However, whether growth of L_{opt} is possible or not may be affected by the particular value of z or λ. Consequently, in the next section our analysis will focus upon the behaviour of $\beta\Omega^*$ as $|\mathcal{O}|$ and L increase, in particular in the regime $|\mathcal{O}|, L \to \infty$, for different choices of z and λ.

5.3. Asymptotic behaviour of the optimal ontology size for regular trees

The second term on the right-hand-side of Eq.(23) is easily evaluated as,

$$\begin{aligned} \omega \log b \sum_{l=0}^{L} l b^l &= \omega \frac{\log b}{(b-1)^2} \left[L b^{L+2} - (L+1) b^{L+1} + b \right] \\ &\simeq \frac{\omega \log b}{(b-1)} L b^{L+1} \quad , \text{ as } L \to \infty \ . \end{aligned} \tag{26}$$

To evaluate the first term on the right-hand-side of Eq.(23) we define $a = b^{\lambda}$ and make use of the Euler-Maclaurin sum formula [21] to obtain (see Appendix A for details),

$$\begin{aligned} \sum_{l=0}^{L} \log\left(1 + z a^l\right) &\simeq \frac{1}{\log a} \left(Li_2(-z) - Li_2(-z a^L) \right) + \frac{1}{2} \log\left(1 + z a^L\right) \\ &+ \frac{1}{2} \log\left(1 + z\right) + \frac{\log a}{12} \left(\frac{z a^L}{1 + z a^L} - \frac{z}{1+z} \right) \ . \end{aligned} \tag{27}$$

From the approximation in Eq.(27) we have,

$$\begin{aligned} \bar{n} &= \frac{1}{\log a} \log\left(\frac{1 + z a^L}{1+z} \right) + \frac{1}{2} \frac{z a^L}{1 + z a^L} + \frac{1}{2} \frac{z}{1+z} \\ &+ \frac{\log a}{12} \left(\frac{z a^L}{(1 + z a^L)^2} - \frac{z}{(1+z)^2} \right) \ . \end{aligned} \tag{28}$$

For convenience of later analysis we can also write Eq.(27) as,

$$\sum_{l=0}^{L} \log\left(1 + z a^l\right) \simeq f_1(z a^L) + f_2(z) \ . \tag{29}$$

with obvious definitions for the functions f_1 and f_2. With this more compact notation we can re-write Eq.(28) as,

$$\bar{n} = z a^L f_1^{'}(z a^L) + z f_2^{'}(z) \ . \tag{30}$$

The approximation developed in Eq.(27) is extremely accurate for the values of z, a and L we are interested in (see Appendix A for details), and thereby

allows us to accurately elucidate the growth behaviour of the optimal ontology size. To proceed, we consider the asymptotic behaviour of the approximation to $\beta\Omega^*$, as $|\mathcal{O}|$ and $L \to \infty$. There are two regimes potentially worth studying in the asymptotic limit $|\mathcal{O}|, L \to \infty$. The first regime is where $\bar{n}$ is fixed in size. In this regime the increasing number of terms available for annotating an object, as $L \to \infty$, are not made use of, *i.e.*, the increased information captured within the larger ontologies is effectively ignored. Consequently, we also find it instructive to consider a second regime where z is fixed, resulting in $\bar{n}$ scaling linearly with L. Detailed analysis of the behaviour of $\beta\Omega^*$ under these two regimes is given below, for both $\lambda < 0$ and $\lambda > 0$.

$\bar{n}$ fixed as $L \to \infty$, $\lambda < 0$:

For $\lambda < 0$ we have $a = b^\lambda < 1$ and so $a^L \to 0$ as $L \to \infty$. Therefore we can expand the right hand side of both Eq.(29) and Eq.(30) in powers of a^L. Doing so, we find,

$$\bar{n} \;=\; z_0 f_2^{'}(z_0) \;\;,\;\; z \;=\; z_0 \;+\; O(a^L) \;\;, \tag{31}$$

$$-\frac{\beta F^*}{|\mathcal{O}|} \;=\; f_1(0) \;+\; f_2(z_0) \;+\; \frac{z_0^2 f_1^{'}(0) f_2^{''}(z_0)}{f_2^{'}(z_0) + z_0 f_2^{''}(z_0)}\, a^L \;+\; O(a^{2L}) \;\;, \tag{32}$$

and so the leading order contributions to $\beta\Omega^*$ take the form,

$$\beta\Omega^* \;\simeq\; \text{Constant} \;-\; |\mathcal{O}| \frac{z_0^2 f_1^{'}(0) f_2^{''}(z_0)}{f_2^{'}(z_0) + z_0 f_2^{''}(z_0)}\, a^L \;+\; \frac{\omega \log b}{(b-1)} L b^{L+1} \quad \text{as } L \to \infty \;\;. \tag{33}$$

The optimal tree depth L_{opt} is determined by minimizing $\beta\Omega^*$ with respect to L. With the constant in Eq.(33) being independent of L we find that, provided $\frac{z_0^2 f_1^{'}(0) f_2^{''}(z_0)}{f_2^{'}(z_0) + z_0 f_2^{''}(z_0)} < 0$, we have,

$$L_{opt} \;\sim\; \frac{\log |\mathcal{O}|}{(1-\lambda)\log b} \quad , \text{ as } |\mathcal{O}| \to \infty \;\;. \tag{34}$$

We show in Appendix B that for $0 > \log a > -6 - 4\sqrt{6} \simeq -15.8$ we indeed have that $\frac{z_0^2 f_1^{'}(0) f_2^{''}(z_0)}{f_2^{'}(z_0) + z_0 f_2^{''}(z_0)} < 0$. So in this regime for $\log a$ Eq.(34) predicts growth of L_{opt}, via a series of transitions, as $|\mathcal{O}|$ is increased, with Eq.(34) giving the global scaling relation between the optimal value L_{opt} and $|\mathcal{O}|$. It is also worth noting that the scaling relation in Eq.(34) predicts that the number of leaf nodes, $b^{L_{opt}}$, or equivalently the number of required classes, grows slower than the number of objects $|\mathcal{O}|$. The restriction to $\log a > -15.8$ is not an onerous one, as it is likely to be well outside the range of values for a that we are interested in (see the end of Appendix B for a discussion of this point).

$\bar{n}$ fixed as $L \to \infty$, $\lambda > 0$:

For $\lambda > 0$ we have $a = b^\lambda > 1$ and so from Eq.(28) we can see that in order to maintain a fixed value of $\bar{n}$ as $L \to \infty$ we must have the scaling $z \sim a^{-L}$. Thus, in general we write $z = \hat{z}_1 a^{-L} + \hat{z}_2 a^{-2L} + O(a^{-3L})$. With $z \geq 0$ then we must

have $\hat{z}_1 \geq 0$ for this decomposition of z to hold over an arbitrary range of L. Substituting this form for z into Eq.(28) we find,

$$\bar{n} \;=\; \hat{z}_1 f_1^{'}(\hat{z}_1) \quad , \quad \hat{z}_2 \;=\; -\frac{\hat{z}_1 f_2^{'}(0)}{f_1^{'}(\hat{z}_1) + \hat{z}_1 f_1^{''}(\hat{z}_1)} \;\; . \tag{35}$$

From this we find,

$$\beta\Omega^* \;\simeq\; -|\mathcal{O}| \left(\text{Constant} \;-\; a^{-L} \hat{z}_1 \hat{z}_2 f_1^{''}(\hat{z}_1) \;+\; O(a^{-2L}) \right) \;+\; \frac{\omega \log b}{(b-1)} L b^{L+1} \;\; . \tag{36}$$

The constant in the expansion above does not depend upon L. In Appendix C we show that $\hat{z}_1 \hat{z}_2 f_1^{''}(\hat{z}_1) < 0$, leading to the immediate conclusion that, as $L \to \infty$, the leading order non-constant term in the asymptotic expansion of the free energy is positive, and so cannot counter-balance the increasing contribution from the complexity penalty term. That is, large tree depths L will never be optimal (in the sense of producing a stationary value of $\beta\Omega^*$) for any value of $|\mathcal{O}|$, irrespective of the values of $\bar{n}$ and a. A natural corollary is that the optimal tree depth, L_{opt}, is then simply the smallest tree that will admit the required value of $\bar{n}$, *i.e.*, $L_{opt} = \bar{n} + 1$, though strictly speaking the leading order asymptotic analysis may not still be valid at such values of L. However, the fact that the asymptotic analysis suggests that using large trees is sub-optimal, irrespective of how many objects we wish to annotate, is surprising.

z fixed as $L \to \infty$, $\lambda < 0$:

With $a < 1$ for $\lambda < 0$ we have $z a^L \to 0$ as $L \to \infty$, and so expanding $\beta\Omega^*$ in powers of a^L gives,

$$\beta\Omega^* \;\simeq\; |\mathcal{O}| \left[\text{Constant} \;-\; z a^L \left(\frac{1}{\log a} + \frac{1}{2} + \frac{\log a}{12} \right) \right] \;+\; \frac{\omega \log b}{(b-1)} L b^{L+1} \;\; . \tag{37}$$

Again, the constant in the expansion of $\beta\Omega^*$ has no dependence upon L. As we have already shown in Appendix B, when $\log a < 0$ we have $\frac{1}{\log a} + \frac{1}{2} + \frac{\log a}{12} < 0$. On setting the derivative (with respect to L) of Eq.(37) to zero we obtain,

$$L_{opt} \;\sim\; \frac{\log |\mathcal{O}|}{(1-\lambda) \log b} \quad , \text{ as } |\mathcal{O}| \to \infty \;\; . \tag{38}$$

Again, we see that it is predicted that the number of classes grows slower than the number of objects $|\mathcal{O}|$.

z fixed as $L \to \infty$, $\lambda > 0$:

For $\lambda > 0$ we have $a = b^{\lambda} > 1$ and so for z fixed we have $z a^L \to \infty$ as $L \to \infty$. Then to obtain the asymptotic behaviour of $\beta\Omega^*$ we observe the following relations and representation for the dilogarithm $Li_2(x)$ [22],

$$\begin{aligned} Li_2(x) \;+\; Li_2(-x) \;&=\; \frac{1}{2} Li_2(x^2) \;\; \forall x \in \mathbb{C} \;\; , \\ Li_2(x) \;&=\; \frac{\pi^2}{3} \;-\; \frac{1}{2} (\log x)^2 \;-\; i\pi \log x \;-\; \sum_{k=1}^{\infty} \frac{1}{k^2 x^k} \quad \forall x \in \mathbb{R}, x \geq 1 \end{aligned} \tag{39}$$

Utilising the results above for $Li_2(x)$ we finally arrive at the leading order behaviour of βF^*,

$$\frac{\beta F^*}{|\mathcal{O}|} \;=\; \text{Constant} \;-\; \frac{L^2}{2}\log a \;-\; L\left(\log z \;+\; \frac{1}{2}\log a\right) \;+\; O\left(a^{-L}\right) \quad . \tag{40}$$

As before, the constant contains only terms that do not depend upon L. Setting $\left.\frac{\partial \beta\Omega^*}{\partial L}\right|_{L=L_{opt}} = 0$, we find the optimal tree depth L_{opt} satisfies the leading order scaling relation,

$$L_{opt} \;\sim\; \frac{\log|\mathcal{O}|}{\log b} \quad \text{as } |\mathcal{O}| \to \infty \;\; . \tag{41}$$

In contrast to the previous scenarios, we now have $|\mathcal{O}|b^{-L_{opt}} \sim$ constant, *i.e.* the number of objects per class (leaf node) is approximately constant (or more correctly, only a weak function of L_{opt}). It is also interesting to observe that with z fixed, the average number of terms used per object, $\bar{n}$, is given by,

$$\bar{n} \;=\; -\frac{z}{|\mathcal{O}|}\frac{\partial}{\partial z}\beta\Omega^* \;\simeq\; \text{Constant} \;+\; L \quad , \text{ as } L \to \infty \;\; . \tag{42}$$

Thus for fixed z we find $\bar{n}$ scaling linearly with L. If we have a measurement of $\bar{n}$ at a particular tree depth L_0, then we can re-express Eq.(42) as,

$$\bar{n}(L) \;=\; \bar{n}(L_0) \;+\; (L \;-\; L_0) \;\; . \tag{43}$$

Summarizing the asymptotic behaviour across the four scenarios we have,

$$\begin{array}{rcccl} L_{opt} \times \log b & \sim & \log|\mathcal{O}|/(1-\lambda) & , & \bar{n} \text{ fixed}, \lambda < 0 \;\; , \\ L_{opt} \times \log b & \sim & \text{Constant} & , & \bar{n} \text{ fixed}, \lambda > 0 \;\; , \\ L_{opt} \times \log b & \sim & \log|\mathcal{O}|/(1-\lambda) & , & z \text{ fixed}, \lambda < 0 \;\; , \\ L_{opt} \times \log b & \sim & \log|\mathcal{O}| & , & z \text{ fixed}, \lambda > 0 \;\; . \end{array} \tag{44}$$

As estimates for $\lambda, \bar{n}$ (and hence z) can already be obtained, the scaling laws (along with their associated amplitudes) provide us with a potential mechanism for assessing whether a tree-like ontology is of optimal size for the given number, $|\mathcal{O}|$, of objects which it is being used to annotate. Thus, we can potentially assess whether an existing ontology should be expanded, or is overly complex for its current usage. Application of these scaling laws is clearly dependent upon their accuracy and the values of λ we are likely to encounter for larger ontologies. These aspects we assess in the next two sections.

5.4. Simulation validation of optimal ontology growth and scaling laws

The possibility of growth in the optimal tree depth L_{opt} arises in three of the four distinct regimes considered above; namely at fixed z for both $\lambda < 0$ and $\lambda > 0$, and also at fixed $\bar{n}$ for $\lambda < 0$. This is borne out by numerical calculations. In Figure 4a we have shown the growth, with $|\mathcal{O}|$, in the tree depth L_{opt}. For fixed $\bar{n}$ with $\lambda < 0$ the fugacity z tends to a finite non-zero value as $L \to \infty$, so essentially we can regard the fixed $\bar{n}, \lambda < 0$ regime as equivalent to the fixed

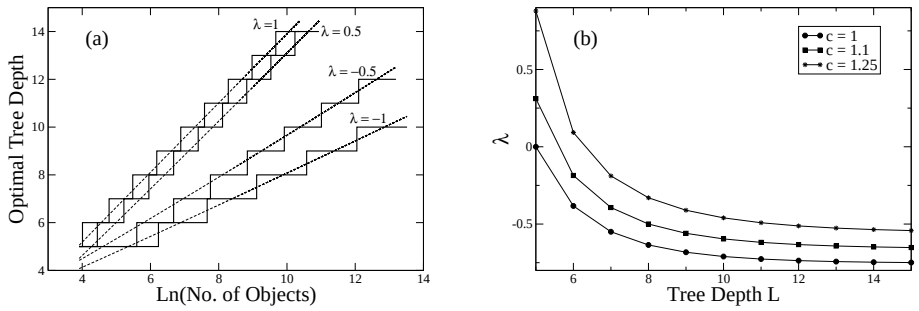

Figure 4: a) Plot of optimal tree depth L_{opt} versus $\ln|\mathcal{O}|$ for various values of λ. The solid lines represent the optimal tree depth L determined by finding the minimum of $\beta\Omega^*$ over integer values of L. The dashed lines represent the broad growth trend as predicted by minimizing the leading order asymptotic contributions to $\beta\Omega^*$. b) Plots of the required values of λ as L increases for a fixed given value of I and $\bar{n}$.

z, $\lambda < 0$ regime. Therefore, we have performed the calculations at a number of different values for λ, but in all cases held the fugacity fixed to achieve a value of $\bar{n} = \frac{25}{7}$ at $L = 5$. All calculations of $\beta\Omega^*$ have used the exact summation form for evaluation of βF^*. The dashed lines shown in Fig.4a correspond to the growth in L_{opt} predicted by by minimizing the leading order asymptotic contributions to the integral approximation of $\beta\Omega^*$. The correspondence between the simulation results and the long term growth trend predicted by the asymptotic analysis is good, confirming the leading order scaling relations given in Eq.(44). The predicted variation with λ (when $\lambda < 0$), in the slope of the scaling relation is also clearly apparent from Fig.4a.

5.5. The expected value of λ for increasing complexity

It is natural to ask which of the above scenarios, $\lambda > 0$ or $\lambda < 0$, is appropriate for a real annotation process? A 'least effort' argument would suggest that an annotator will attempt to limit the effort expended on annotating an object, and so I, the information retrieved per object, will most likely be fixed or only a weakly increasing function of $|\mathcal{O}|$. At fixed z simple inspection of Eq.(25) shows that I is an increasing function of tree depth L_{opt}, and hence $|\mathcal{O}|$, for all values of λ. However, more detailed analysis of Eq.(25) (in particular applying the Ratio test for series convergence, when $\lambda < 0$) reveals that as $L \to \infty$ we have (at fixed z),

$$\begin{aligned} I &\sim \frac{\log b}{2} L^2 + O(L) \quad , \quad \lambda > 0 \ , \\ I &\to \text{Constant} \quad , \quad \lambda < 0 \ . \end{aligned} \tag{45}$$

Therefore, we associate $\lambda < 0$ with a fixed, or only weakly growing value of I. This idea is corroborated if we consider the more intuitive scenario of a

fixed value of $\bar{n}$ (as opposed to fixed fugacity z) and a fixed value of I, and determine the required value of λ as the tree depth is L is increased. Figure 4b shows plots, against L, of the value of λ required. The value of $\bar{n}$ is held fixed at $\bar{n} = 25/7$ and we have fixed $I(\lambda, L, \bar{n}) = c \times I(\lambda = 0, L = 5, \bar{n})$ for different values of c. The value of λ is then obtained by simultaneous solution of Eq.(24) and Eq.(25). Irrespective of the value of c we can see that increasing L at fixed I and $\bar{n}$ leads to $\lambda < 0$. Consequently, if the information or effort expended per object by an annotator is limited we expect small negative values of λ to be the norm. Furthermore, we note that the amount of annotation information retrieved per object essentially determines the specificity with which the different classes of objects are discriminated. Thus, as data sets of increasing size will potentially sample an increasing number of the classes $\mathcal{P}$ present in the population, a fixed value of I may not be sufficient to discriminate them, *i.e.*, the annotation data set will not be of adequate quality. Therefore, the value of λ estimated from a real data set provides a topology derived metric to assess the quality of annotation and annotators, with a negative estimate for λ giving possible cause for concern.

6. Discussion and Conclusions

Although the real-world ontologies used in annotating data sets may be fixed in form, the annotation process, by nature of it often being a manual process, is subject to variation. We have used statistical mechanics to construct a formalism of that variability in the annotation process. The formalism developed has allowed us to understand both the patterns seen in real data sets [5], and to suggest measures of the ontology structure itself. This has been done by using a simple lattice-gas model of the term selection process combined with information theoretic concepts. Although measures of ontology structures [9] already exist these do not tend to incorporate the effects of the expected variability in the annotation process. Likewise, many studies have previously combined information theoretic concepts with ontologies, but these works have largely not focused on assessing the underlying ontology structure. Instead studies have used information theoretic concepts in assessing the similarity of individual terms within an ontology [23, 24, 25], assessing the similarity between annotated objects [26], assessing the similarity between two annotation data sets [27], assessing the similarity between users [28], or in empirical studies of the growth in tag frequencies [29]. Similarly, studies of annotation and tag statistics that relate to the underlying ontology structure exist within the physics literature [14, 15], but these have not been based upon a Hamiltonian model of the annotation process. By constructing a Hamiltonian model of the annotations we have built upon the work of Palla *et al.* [14] and Tibély *et al.* [15], and been able to construct, in a principled fashion, statistical mechanical measures of the ontology itself.

Having developed a simple lattice-gas model of the annotation process the statistical mechanical formalism has enabled us to progress further and to perform hypothetical experiments over an ensemble of different ontology structures, and in doing so gain new insight where it is not readily possible to perform real

experiments. Our analysis of an ensemble of different ontology structures has focused on regular trees, though despite this restriction we still expect the high-level conclusions drawn in Sections 5.2, 5.4, and 5.5 to be more universally valid. Firstly, our detailed analysis of $\beta\Omega^*$ identifies a natural or optimal ontology size, and associated growth scaling law, given the number of objects $|\mathcal{O}|$ to be annotated. The scaling laws derived provide us with 'rules-of-thumb' to decide when an ontology should be expanded to match the needs of the data sets to which it is being applied. Secondly, our analysis in Section 5.5 reveals a natural tendency towards $\lambda < 0$, as more complex ontologies are used to annotate larger collections of objects. This suggests that values of λ appropriate to real-world annotation data sets will be typically small in magnitude, possibly even negative. This is borne out by the power-law exponents observed by Kalankesh *et al.* [5]. Small values of λ would imply that annotators are typically recovering less information from the ontology than if they selected terms uniformly at random. As we equate effort expended with the information retrieved, a negative value of λ suggests the effort being expended during the annotation is below that appropriate to the ontology complexity. This may be due to the structure of the ontology being more complex than is necessary to discriminate between the classes of objects being annotated, the annotation effort expended not being sufficient, or simply that there is currently insufficient evidence available to the annotators to discriminate between certain objects. With the statistical mechanical formalism we can potentially extend the analysis presented here to construct measures capable of distinguishing these different possibilities.

The statistical mechanical analysis we have presented here is based upon a relatively simple Hamiltonian. The richness of the behaviour we observe in the statistical mechanical model is more a consequence of the inhomogeneous field induced by the ontology topology. However, the statistical mechanical analysis we have presented here is far from complete. The advantage of having expressed the annotation process in terms of a statistical mechanical formalism is that we can easily extend our analysis to obtain quantitative results for other scenarios or other Hamiltonians. For example, to further reflect more realistic annotation patterns the analysis could be extended to take into account the non-independence of term usage. Indeed, term co-occurrence frequencies have been studied by Tibely *et al.* [15], with terms that are closer (as measured by path length on the ontology DAG) occurring together more frequently than those further apart. Incorporating tag co-occurrence probabilities would turn our lattice-gas model from a non-interacting model to an interacting model. Alternatively, our statistical mechanical analysis can potentially be extended to account for more realistic aspects and nuances of the annotation process, *e.g.*, mis-selection of terms that are not appropriate to the object being annotated.

Appendix A.

Applying the Euler-Maclaurin sum formula to $\sum_{l=0}^{L}\log\left(1+zb^{l\lambda}\right)$ gives, to lowest order of expansion [21],

$$\begin{aligned}
&\sum_{l=0}^{L}\log\left(1+zb^{l\lambda}\right) \;=\; \sum_{l=0}^{L}\log\left(1+za^{l}\right)\\
=\;& L\int_0^1 ds\log\left(1+za^{sL}\right) \;+\; \frac{1}{2}\log\left(1+za^{L}\right) \;+\; \frac{1}{2}\log\left(1+z\right)\\
+\;& \frac{1}{2}B_2\log a\left(\frac{za^{L}}{1+za^{L}} \;-\; \frac{z}{1+z}\right) \;-\; \frac{(\log a)^2}{2}\int_0^L dsB_2\left(s-\lfloor s\rfloor\right)\frac{za^{s}}{(1+za^{s})^2} \;\;,
\end{aligned} \tag{A.1}$$

where B_2 is the second Bernoulli number, $B_2(x)=x^2-x+\frac{1}{6}$ is the second Bernoulli polynomial, and $a=b^{\lambda}$. The first integral on the right-hand-side of Eq.(A.1) can be expressed in terms of the dilogarithm function, $Li_2(x)$ [30, 22], as,

$$L\int_0^1 ds\log\left(1+za^{sL}\right) \;=\; \frac{1}{\log a}\left(Li_2(-z) \;-\; Li_2(-za^{L})\right) \;\;. \tag{A.2}$$

With $\frac{za^{s}}{(1+za^{s})^2}>0\;\forall s\in[0,L]$ we can bound the last term in Eq.(A.1) simply by replacing $B_2(s-\lfloor s\rfloor)$ by its minimum and maximum values in the integrand ($-\frac{1}{12}$ and $\frac{1}{6}$, respectively). This yields bounds for $\sum_{l=0}^{L}\log\left(1+za^{l}\right)$ as,

$$\begin{aligned}
&\frac{1}{2}\log\left(1+za^{L}\right) \;+\; \frac{1}{2}\log\left(1+z\right) \;\leq\; \sum_{l=0}^{L}\log\left(1+za^{l}\right)-L\int_0^1 ds\log\left(1+za^{sL}\right)\\
\leq\;& \frac{1}{2}\log\left(1+za^{L}\right) \;+\; \frac{1}{2}\log\left(1+z\right) \;+\; \frac{\log a}{8}\frac{za^{L}}{1+za^{L}} \;-\; \frac{\log a}{8}\frac{z}{1+z} \;\;.
\end{aligned} \tag{A.3}$$

To arrive at a final approximation to $\sum_{l=0}^{L}\log\left(1+za^{l}\right)$ that is convenient for further analysis we simply take some point between the two bounds in Eq.(A.3). By replacing $B_2(s-\lfloor s\rfloor)$ by its average value of zero, in the last integral in Eq.(A.1), we arrive at the approximation,

$$\begin{aligned}
\sum_{l=0}^{L}\log\left(1+za^{l}\right) \;\;\simeq\;\; & \frac{1}{\log a}\left(Li_2(-z) \;-\; Li_2(-za^{L})\right) \;+\; \frac{1}{2}\log\left(1+za^{L}\right)\\
&+\;\; \frac{1}{2}\log\left(1+z\right) \;+\; \frac{\log a}{12}\left(\frac{za^{L}}{1+za^{L}} \;-\; \frac{z}{1+z}\right) \;\;.
\end{aligned} \tag{A.4}$$

We should note that the approximation in Eq.(A.4) is simply equivalent to dropping the remainder term in the representation of the Euler-Maclaurin sum formula we have used.

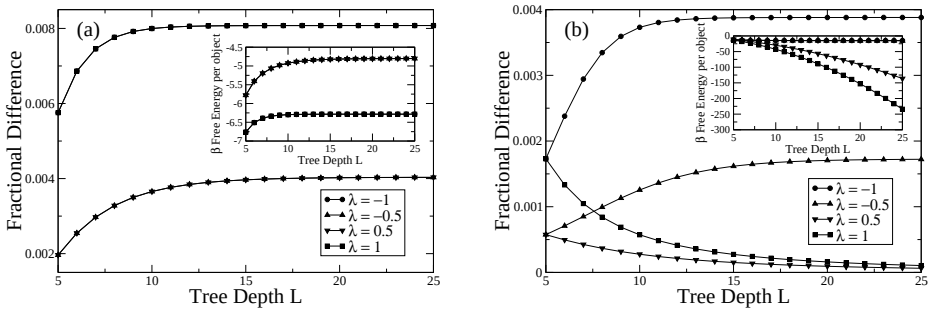

Figure A.5: Plots versus tree depth L, at different values of λ, of the fractional difference between the exact free energy for a regular tree and the free energy calculated using the approximation Eq.(A.4). The inset shows the exact free energy values. In all plots we have set the number of objects $|\mathcal{O}| = 1000$. In a) the number of annotation terms per object, $\bar{n}$, has been fixed at $\bar{n} = 25/7$, whilst in b) the fugacity z has been fixed to achieve a value of $\bar{n} = 25/7$ when $L = 5$.

The approximation developed in Eq.(A.4) is in fact extremely accurate for the values of z, a and L we are interested in. Figure A.5 shows the fractional difference between the exact evaluation of the sum on the left-hand-side of Eq.(A.4) and the approximation on the right-hand-side of Eq.(A.4). Fig.A.5a shows the fractional difference for different values of L and λ at a fixed values of $\bar{n}$ and suggests the approximation in Eq.(A.4) is typically accurate to within 1 part in 100 for fixed $\bar{n}$. For the plots in Fig.A.5a the same value for the fugacity z has been used in both free energy calculations. Namely, we have used the value of z required to achieve a value of $\bar{n} = 25/7$, as determined from the exact summation relation in Eq.(24). Similarly, in Fig.A.5b shows the fractional difference between the exact summation in Eq.(A.4) and the integral based approximation, at different values of L and λ but for a fixed value of z. Again, the same fixed value for the fugacity z has been used in both free energy calculations, namely that required to achieve a value of $\bar{n} = 25/7$ when $L = 5$, as determined from exact summation relation in Eq.(24). Fig.A.5b would suggest that for fixed z the approximation in Eq.(A.4) is again accurate to 2 parts in 1000.

Appendix B.

To observe growth in L_{opt} as $|\mathcal{O}| \to \infty$, at fixed $\bar{n}$ with $\lambda < 0$, we require,

$$z_0^2 \frac{f_1^{'}(0) f_2^{''}(z_0)}{f_2^{'}(z_0) + z_0 f_2^{''}(z_0)} \; < \; 0 \; . \tag{B.1}$$

We start by analysing the behaviour of the numerator of Eq.(B.1). From Eq.(27)

and Eq.(29) we have that $f_1(x)$ is defined as,

$$f_1(x) \;=\; -\frac{1}{\log a} Li_2(-x) \;+\; \frac{1}{2}\log(1+x) \;+\; \frac{\log a}{12}\frac{x}{1+x} \;\;, \tag{B.2}$$

from which we have,

$$f_1^{'}(0) \;=\; \frac{1}{\log a} \;+\; \frac{1}{2} \;+\; \frac{\log a}{12} \;\;. \tag{B.3}$$

We wish to determine the behaviour of $f_1^{'}(0)$ for $a \in \mathbb{R}$ and $0 < a < 1$. To do so we seek the roots of $f_1^{'}(0) = 0$ as a function of $y = \log a < 0$. That is we solve,

$$\begin{aligned} & \frac{1}{y} \;+\; \frac{1}{2} \;+\; \frac{y}{12} = 0 \quad , \;\; y \;=\; \log a \\ \Rightarrow \quad & y^2 \;+\; 6y \;+\; 12 \;=\; 0 \;\;. \end{aligned} \tag{B.4}$$

As the discriminant of this quadratic equation is -12, there are no real roots. Therefore $f_1^{'}(0)$ maintains the same sign for all real values of $y < 0$. It is a simple matter to confirm that $f_1^{'}(0) = -\frac{7}{12}$ at $y = -1$ and so $f_1^{'}(0) < 0$ for all $\lambda < 0$ with $b > 1$.

From Eq.(27) and Eq.(29) we have that $f_2(x)$ is defined as,

$$f_2(x) \;=\; \frac{1}{\log a} Li_2(-x) \;+\; \frac{1}{2}\log(1+x) \;-\; \frac{\log a}{12}\frac{x}{1+x} \;\;, \tag{B.5}$$

from which we find,

$$f_2^{''}(x) \quad = \quad \frac{1}{\log a}\left(\frac{1}{x^2}\log(1+x) \;-\; \frac{1}{x(1+x)}\right) \;-\; \frac{1}{2}\frac{1}{(1+x)^2} \;+\; \frac{\log a}{6}\frac{1}{(1+x)^3} \;\;. \tag{B.6}$$

With $\log a < 0$ it is clear that if $\log(1+x) > x/(1+x) \;, \forall x > 0$ then we will have $f_2^{''}(x) < 0 \;, \forall x > 0$. At $x = 0$ we have $\log(1+x) = x/(1+x)$. We also have,

$$\begin{aligned} & \frac{d\log(1+x)}{dx} \;=\; \frac{1}{1+x} \quad , \quad \frac{d[x/(1+x)]}{dx} \;=\; \frac{1}{(1+x)^2} \\ \Rightarrow \quad & \frac{d\log(1+x)}{dx} \;\geq\; \frac{d[x/(1+x)]}{dx} \quad \forall x \geq 0 \;\;, \\ \Rightarrow \quad & \log(1+x) \;\geq\; \frac{x}{1+x} \quad \forall x \geq 0 \;\;, \end{aligned} \tag{B.7}$$

With $z_0^2 > 0$ we find that the numerator of Eq.(B.1) is negative.

Now considering the behaviour of the denominator of Eq.(B.1) we have,

$$f_2^{'}(z_0) + z_0 f_2^{''}(z_0) \;=\; \frac{d\bar{n}(z_0)}{dz_0} \;\;. \tag{B.8}$$

For an exactly calculated partition function we might expect $\bar{n}(z_0)$ to be monotonically increasing in z_0. However, as we are utilising an integral approximation to the exact expression for $\beta\Omega^*$ given in Eq.(23), we continue our analysis with $d\bar{n}(z_0)/dz_0$ as given above. Explicitly, we find,

$$f_2^{'}(x) + x f_2^{''}(x) \ = \ -\frac{\left[(12 - 6y + y^2) \ + \ x(24 - 6y - y^2) \ + \ 12x^2\right]}{12y(1+x)^3} \ . \quad \text{(B.9)}$$

where again $y = \log a < 0$. Examination of the discriminant (with respect to x) of the numerator in the right-hand-side of Eq.(B.9) reveals that in order for the right-hand-side of Eq.(B.9) to be positive requires $\log a < -6 - 4\sqrt{6} \simeq -15.8$. Therefore, overall we find that the left-hand-side of Eq.(B.1) is negative for $0 > \log a > -6 - 4\sqrt{6}$. This threshold value of a we regard as somewhat extreme and outside the range that we are likely to be interested in - a value of $\log a = -15.8$ corresponds to a reduction by a factor of more than 10^7 in the probability of selecting a term, in going from one layer of the tree to the next layer. The possibility that $d\bar{n}(z)/dz$ may be negative for some z in the limit $L \to \infty$ is most likely due to a breakdown in the accuracy of the integral approximation in Eq.(A.1). Indeed, from Eq.(24) we see that the exact expression for $d\bar{n}(z)/dz$ is given by,

$$\frac{d\bar{n}}{dz} \ = \ \frac{1}{z}\sum_{l=0}^{L} \frac{1}{(1+za^l)^2} \ , \quad \text{(B.10)}$$

which is clearly positive for all λ and for all $z > 0, a > 0$.

Appendix C.

For fixed $\bar{n}$ and $\lambda > 0$ the leading order contribution to $\beta\Omega^*$ from the free energy, that is dependent upon L, is given by,

$$|\mathcal{O}|a^{-L}\hat{z}_1\hat{z}_2 f_1^{''}(\hat{z}_1) \ . \quad \text{(C.1)}$$

Using the relations in Eq.(35) we can re-write $\hat{z}_1\hat{z}_2 f_1^{''}(\hat{z}_1)$ as,

$$\hat{z}_1\hat{z}_2 f_1^{''}(\hat{z}_1) \ = \ -\frac{\hat{z}_1^2 f_1^{''}(\hat{z}_1) f_2^{'}(0)}{f_1^{'}(\hat{z}_1) + \hat{z}_1 f_1^{''}(\hat{z}_1)} \ . \quad \text{(C.2)}$$

From the definition of $f_1(x)$ given in Eq.(B.2) we find that,

$$f_1^{'}(\hat{z}_1) + \hat{z}_1 f_1^{''}(\hat{z}_1) \ = \ \frac{1}{\log a}\frac{1}{1+\hat{z}_1} \ + \ \frac{1}{2}\frac{1}{(1+\hat{z}_1)^2} \ + \ \frac{\log a}{6}\frac{2+\hat{z}_1}{(1+\hat{z}_1)^3} \ . \quad \text{(C.3)}$$

With $\hat{z}_1 \geq 0$ and $a > 1$ we find that the denominator in Eq.(C.2) is positive. Of the three factors contributing to the numerator in Eq.(C.2), $\hat{z}_1^2$ is clearly positive. The form for $f_2(x)$ is already defined in Eq.(B.5), from which we find,

$$f_2^{'}(0) \ = \ -\frac{1}{\log a} \ + \ \frac{1}{2} \ - \ \frac{\log a}{12} \ . \quad \text{(C.4)}$$

As with Eq.(B.3) we find that $f_2^{'}(0) = 0$ has no real roots and so has the same sign for all $\log a > 0$. At $\log a = 1$ we find $f_2^{'}(0) = -\frac{7}{12}$ and so conclude that $f_2^{'}(0) < 0\ , \forall \log a > 0$.
For the remaining factor in the numerator of Eq.(C.2) we have,

$$f_1^{''}(x) \;=\; \frac{1}{\log a}\left(-\frac{1}{x^2}\log(1+x) \;+\; \frac{1}{x(1+x)}\right) \;-\; \frac{1}{2}\frac{1}{(1+x)^2} \;-\; \frac{\log a}{6}\frac{1}{(1+x)^3} \quad . \tag{C.5}$$

With $\log a > 0$ it is clear that we will have $f_1^{''}(x) < 0\ , \forall x > 0$ if $\log(1+x) > x/(1+x)\ , \forall x > 0$. This latter relationship has already been proved in Eq.(B.7), from which we can ultimately conclude $f_1^{''}(\hat{z}_1) \leq 0\ \ \forall \hat{z}_1 \geq 0$. Combining the three inequalities above relating to the numerator and denominator of Eq.(C.2) we finally conclude that,

$$\hat{z}_1\hat{z}_2 f_1^{''}(\hat{z}_1) \;=\; -\frac{\hat{z}_1^2 f_1^{''}(\hat{z}_1) f_2^{'}(0)}{f_1^{'}(\hat{z}_1) + \hat{z}_1 f_1^{''}(\hat{z}_1)} \;\leq\; 0 \;\;,\;\; \forall \hat{z}_1 \geq 0 \;\; . \tag{C.6}$$

References

[1] A.G. Taylor and D.N. Joudrey. *The Organization of Information*. Libraries Unlimited, Westport, CT, 3rd edition, 2008.

[2] J. Davies, A. Duke, and Kiryakov A. Semantic search. In J. Davies and A. Göker, editors, *Information Retrieval: Searching in the 21st Century*. John Wiley and sons, Chichester, 2009.

[3] I. Peters and W.G. Stock. Folksonomy and information retrieval. *Proceedings of the American Society for Information Science and Technology*, 44(1):1–28, 2007.

[4] T. Gruber. Ontology. In L. Liu and M. Tamer Özsu, editors, *Encyclopedia of Database Systems*, London, 2008. Springer-Verlag.

[5] L. R. Kalankesh, R. Stevens, and A. Brass. The language of gene ontology: a zipf's law analysis. *BMC Bioinformatics*, 13:127, 2012.

[6] M. Ashburner, C.A. Ball, J.A. Blake, D. Botstein, H. Butler, and J.M. Cherry et. al. Gene ontology: tool for the unification of biology. *Nature Genetics*, 25(1):25–29, 2000.

[7] G.K. Zipf. *The Psychobiology of Language: An Introduction to Dynamic Philology*. Routledge, London, 1936.

[8] G.K. Zipf. *Human Behaviour and the Principle of Least Effort*. Addison-Wesley, Cambridge, MA, 1949.

[9] L. Yao, A. Divoli, I. Mayzus, J. A. Evans, and A. Rzhetsky. Benchmarking ontologies: Bigger or better? *PLoS Computational Biology*, 7:e1001055, 2011. doi:10.1371/journal.pcbi.1001055.

[10] R. Ferrer i Cancho and Solé R.V. Least effort and the origins of scaling in human language. *Proceedings of the National Academy of Sciences*, 100:788–791, 2003.

[11] R. Ferrer i Cancho. Zipf's law from a communicative phase transition. *European Physical Journal B*, 47:449–457, 2005.

[12] R. Ferrer i Cancho and A. Diaz-Guilera. The global minimum of the communicative energy of natural communication systems. *Journal of Statistical Mechanics*, 2007. doi:10.1088/1742-5468/2007/06/P06009.

[13] K. Kosmidis, A. Kalampokis, and P. Argyrakis. Statistical mechanical approach to human language. *Physica A*, 366:495–502, 2006.

[14] G. Palla, I.J. Farkas, P. Pollner, I. Derényi, and T. Vicsek. Fundamental statistical features and self-similar properties of tagged networks. *New Journal of Physics*, 10:123026, 2008.

[15] G. Tibély, P. Pollner, T. Vicsek, and G. Palla. Ontologies and tag-statistics. *New Journal of Physics*, 14:053009, 2012.

[16] J. Bang-Jensen and G.Z. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer-Verlag, London, 2nd edition, 2009.

[17] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 3rd edition, 2009.

[18] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, 1991.

[19] H.W. Watson and F. Galton. On the probability of the extinction of families. *Journal of the Anthropological Institute of Great Britain*, 4:138–144, 1875.

[20] D.G. Kendall. Branching processes since 1873. *Journal of the London Mathematical Society*, s1-41(1):385–406, 1966.

[21] E.T. Whittaker and G. N. Watson. *A Course in Modern Analysis*. Cambridge University Press, Cambridge, England, 4th edition, 1990.

[22] A.N. Kirillov. Dilogarithm identities. *Progress of Theoretical Physics Supplement*, 118:61–142, 1995.

[23] P. Resnik. Using information content to evaluate semantic similarity in a taxonomy. In *IJCAI1995, Proceedings of the 14th international joint conference on artificial intelligence*, pages 448–453, San Francisco, CA, 1995. Morgan Kaufmann Publishers Inc.

[24] P. Resnik. Semantic similarity in a taxonomy: an information-based measure and its application to problems of ambiguity in natural language. *Journal of Artificial Intelligence Research*, 11:95–130, 1999.

[25] J. Calmet and A. Daemi. From entropy to ontology. In *Cybernetics and Systems 2004 - AT2AI-4: From agent theory to agent implementation*, 2004.

[26] Y. Tao, L. Sam, J. Lee, C. Friedman, and Y. Lussier. Information theory applied to the sparse gene ontology annotation network to predict novel gene function. *Bioinformatics*, 23:i529–i538, 2007.

[27] G. Alterovitz, M. Xiang, and M.F. Ramoni. An information theoretic framework for ontology-based bioinformatics. In *ITA2007, Proceedings of Information Theory and Applications Workshop*, pages 16–19, 2007. doi=10.1109/ITA.2007.4357555.

[28] W. Koh and L. Mui. An information theoretic approach to ontology-based interest matching. In A. Maedche, S. Staab, C. Nedellec, and Hovy E., editors, *Proceedings of the Second Workshop on Ontology Learning OL'2001*, RWTH Aachen University, 2001. CEUR-WS. available from http://ceur-ws.org/Vol-38/koh.pdf.

[29] E.H. Chi and T. Mytkowicz. Understanding the efficiency of social tagging systems using information theory. In *HyperText08, Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*, pages 81–88, New York, 2008. ACM.

[30] D. Zagier. The dilogarithm function. In P. Cartier, P. Moussa, B. Julia, and P. Vanhove, editors, *Frontiers in Number Theory, Physics, and Geometry II*, pages 3–65. Springer Berlin Heidelberg, 2007.